marvin

The Modern Researcher's Guide

Building a Customer Context Layer for AI Tools



Table of Contents

Why your team's AI tools need customer context 03

Three reasons why shared context can't wait 05

Anatomy of a customer context layer 08

Who should own the context layer 10

Steps to build the context layer 12

What this looks like with Marvin 18

◆ Introduction

Why your team's AI tools need customer context

Your AI tools have never met your customers. Yet, when working in Claude or ChatGPT, it's very likely that your tool won't hesitate to tell you what customers want, and with total confidence. But the reality is, AI is only as smart as the information you give it, which is why "context" has become the buzzword of 2026. Unfortunately, in most organizations, the customer knowledge that AI needs is scattered, sitting in transcripts, support tickets, survey responses, and interview recordings it can't reach. Without that evidence, it runs at a fraction of its potential.

In our State of Modern Research 2026 report, a survey of 309 researchers, research consumers, and research leaders, **59%** reported using general-purpose AI tools. But **nearly half** of these teams have no shared source of customer knowledge.

What evidence will AI use instead? One AI assistant has the PDF a team uploaded last spring, while another is looking at only a few support tickets pasted into a chat. A third has nothing at all and fills in the gaps with assumptions the AI made up itself.

Product leaders have a name for this siloed approach: **single-player** AI. Arnab Bose, chief product officer at Asana, describes it as AI tools that "assist one person, in one conversation, disconnected from the wider business." He argues that the way out is building the infrastructure to provide context. "In a landscape where every organisation has access to roughly the same underlying models, the quality of the context you provide — and the infrastructure you build to support it — may turn out to be the most important competitive variable of all."

That's the context gap: the distance between how widely teams use AI and how little shared context their AI can reach, customer context included.

The alternative is **multiplayer AI**, where people and AI tools work from shared plans, workflows, and context. Customer knowledge is one of the most valuable parts of that shared context, and it needs infrastructure of its own. That infrastructure is a customer context layer.

The good news is: many teams already have the foundation to create this shared resource. For years, researchers have built systems where customer data flows in, gets organized and analyzed, and reaches the teams making decisions. **Most companies don't treat those systems as part of their AI strategy yet.** The ones that do are already reaping benefits. In our survey, respondents with a shared repository as a context layer are more than **three times as likely to see significant AI time savings (51% vs. 16%)**.

A customer context layer is a single, trusted source of everything your customers have told you and keep telling you, available to every team and every AI tool.

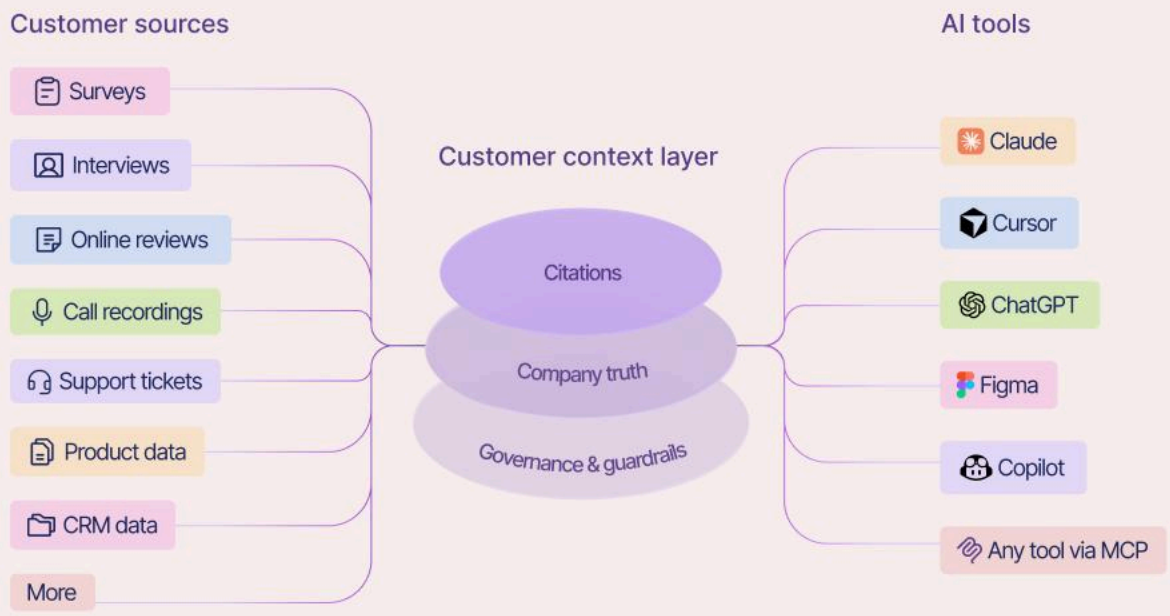

Connecting the context layer to each tool is no longer a challenge. MCP (The Model Context Protocol) is an open standard that lets AI platforms such as Claude, ChatGPT, and Copilot ask outside sources for context. But MCP needs help too. When you set up the context layer, you need to organize the data, manage who can see it, and protect personal customer information; or use a platform that lets you control all three.

But why do you need to prioritize a context layer now? Let's look at three reasons why your team urgently needs to build that shared resource.

Three reasons why shared context can't wait

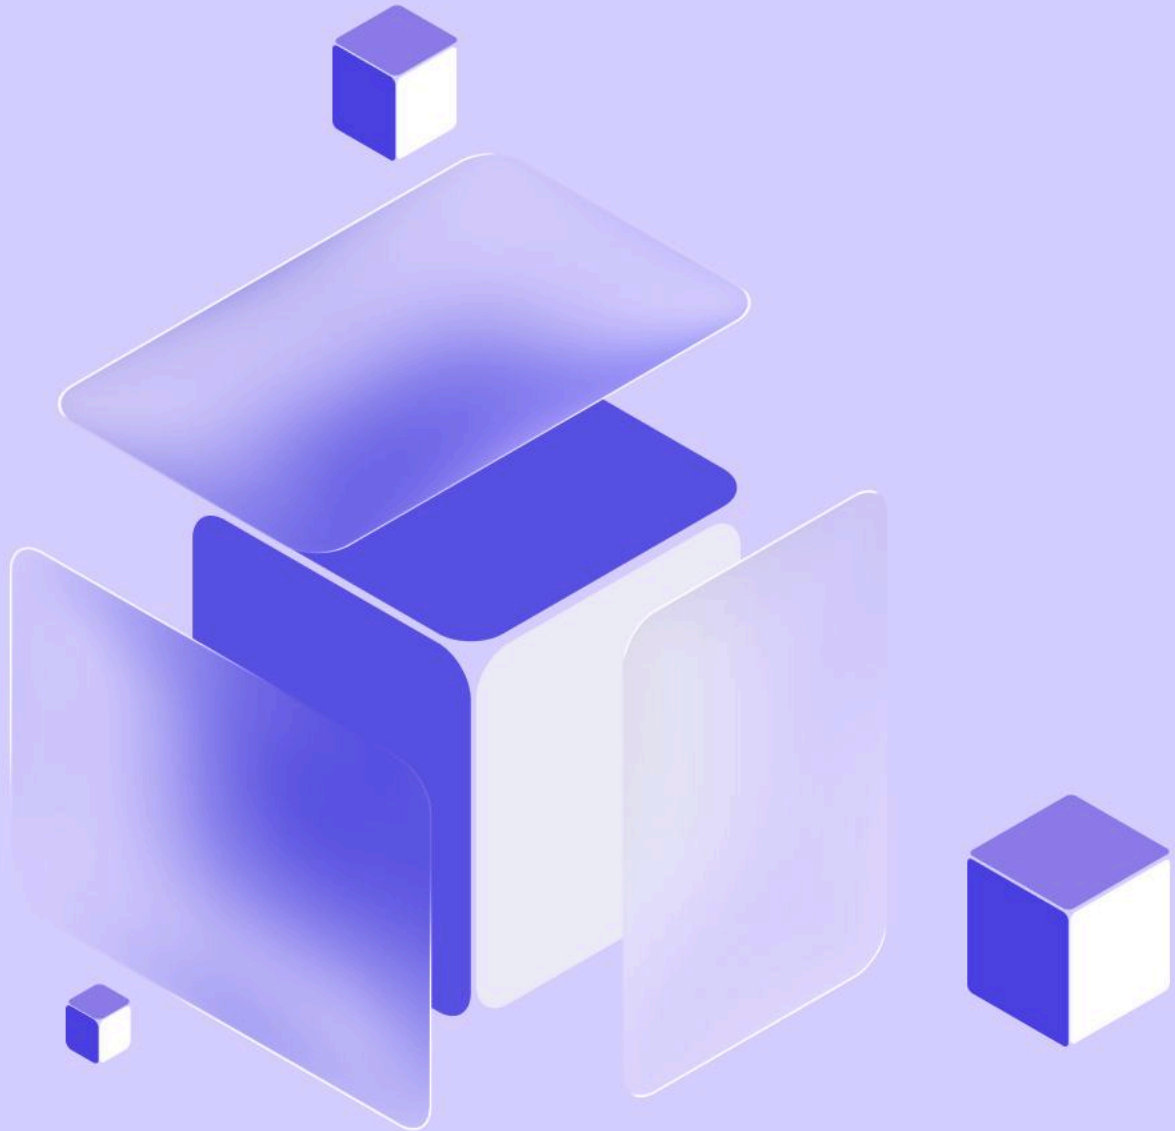

◆ Why now

Three reasons why shared context can't wait

1. Every AI tool tells a different story about your customers

AI use varies a lot person by person, and most of what people learn stays siloed with them. The most advanced AI users are the exception. Microsoft's 2026 Work Trend Index calls them Frontier Professionals, and they were more likely than everyone else to say their teams "share AI tips, new agents, learnings, and mistakes" **(61% vs. 36%)**.

That's multiplayer AI in practice. But most people still work single-player, with each tool only knowing what its user gave it. When the source material differs significantly, so do the answers. A product manager and a designer can ask about the same customer and walk away with conflicting conclusions.

2. Messy customer data leads to unreliable AI answers

In February 2025, Gartner predicted: "Through 2026, organizations will abandon 60% of AI projects unsupported by AI-ready data." AI-ready data is organized, labeled, and complete enough for AI to draw the right conclusions, or, in Gartner's words, "representative of the use case." Customer knowledge is hard to make AI-ready, since most of it lives in recordings, open-ended survey answers, and interview notes.

Pointing AI at all of it doesn't fix the problem. Search can be selective, and it can't tell you which insights are representative, current, or true. On our How Modern Researchers Are Using MCP webinar, Kaleb Loosbrock, founder of the AIxUXR community, named the requirement: "Clean, structured data is core to facilitating and wrangling AI to get quality outputs."

3. If research can't reach AI tools, it can't shape decisions

Researchers have spent years making insights easier for stakeholders to reach, because the easier research is to find, the more it impacts the business. Now, more of those stakeholders do their work and ask their questions inside AI tools, so that's where research needs to show up. In August 2025, Gartner predicted that "40% of enterprise applications will be integrated with task-specific AI agents by the end of 2026, up from less than 5% today." When an agent can't reach your research, it answers without it, and decisions about customers get made without the evidence your team gathered.

What it takes to close the context gap

All three reasons point to the same two needs:

Bring every source together. Customer knowledge is spread across research studies, sales calls, support tickets, reviews, and more. In our survey, **61% of respondents** reported scattered research data, conflicting evidence across sources and departments, or no source of truth at all. An AI tool that reaches only one of these sources gives a partial answer and presents it as complete.

Keep every answer tied to its source. A claim about customers is only as strong as the evidence behind it. Traceability has always been a core research principle, and it matters even more when AI does the answering. When every answer links back to who said it, when it was said, and in what context, team members can check the validity before they act on it, and the research team stays connected to the decisions it helps improve. Pasted snippets lose that trail. Without it, a comment from one vocal customer looks the same as a theme heard across fifty interviews.

That takes a layer built for customer knowledge.

Anatomy of a customer context layer

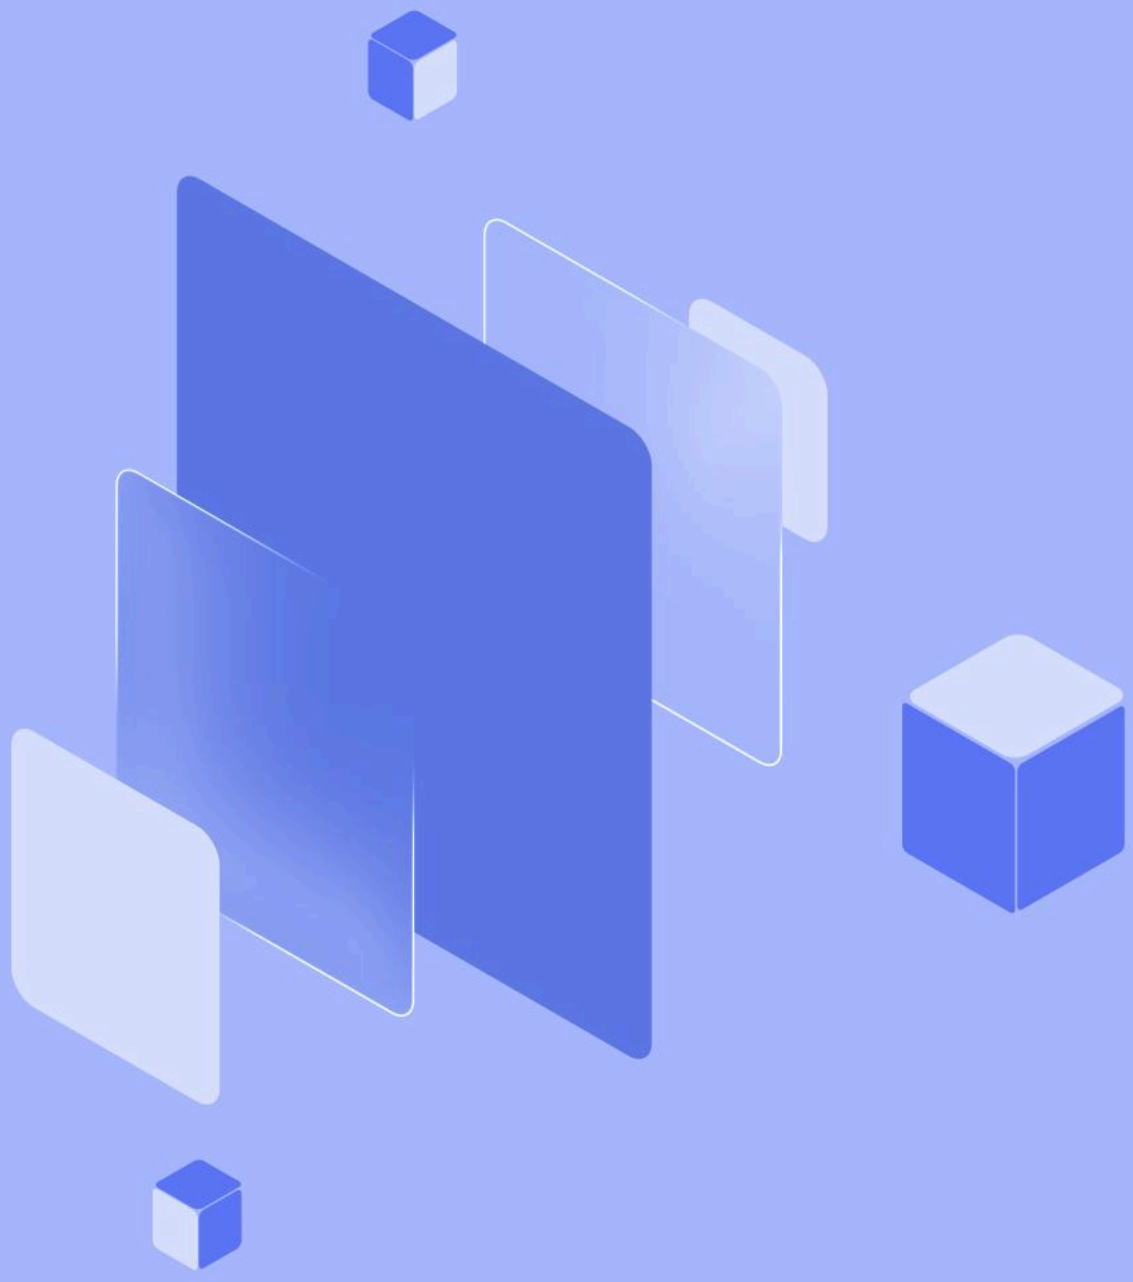

◆ The solution

Anatomy of a customer context layer

The goal is to build a context layer that takes everything your customers tell you and keeps it structured, cited, and governed by your company's data sharing rules. Every team and every AI tool then draws on the same trusted source.

Think of it as one customer brain for many players: researchers, product teams, marketers, and their AI tools.

COMPLETE	Interviews, surveys, support tickets, sales calls, reviews, and every ongoing feedback channel feed into one place.
TRUSTWORTHY	Every AI answer about customers is one you can trust, cited to its source and grounded in your company's terminology, product facts, and decisions. Governance travels with it, so permissions and privacy rules apply to every result.
CONNECTED	It reaches people and AI tools where they already work, from Slack and Figma to Claude, Copilot, and internal agents, with MCP as the common route in. AI tools ask it for customer evidence before they answer.
LIVING	The layer keeps updating as customers keep talking, through everyday channels and through ongoing studies. It learns as it ingests more information and grows more valuable over time.

Together, the four properties turn **scattered customer data into shared context.**

Who should own the context layer



Ownership

Who should own the context layer

Research and insights teams are the natural owners of the layer, because understanding customers is the heart of their work. Researchers know which evidence is solid and which is anecdotal, how to organize findings so they hold up, and when a pattern is strong enough to act on. They also know how to handle customer data responsibly, from consent to privacy. The layer needs someone to make exactly those calls about what goes in, how it's organized, and which findings count as answers. Companies without a dedicated research team still need a named owner for those calls, often a product or design leader or a product operations team, working with the people who manage each source.

When researchers run a shared system for customer knowledge, research carries more weight. In our State of Modern Research data, **85%** of organizations with an actively used repository reference research in major decisions often or almost always, compared with **60%** of those without one. They're also **2.5 times** as likely to track both the decisions research informs and the outcomes of those decisions **(87% vs. 35%)**.

Owning the layer also means deciding who sees what. There's no single right setting, and each choice carries a trade-off.

Eniola Abioye, founder of UX Outloud and a former UX researcher at Meta, wants people to reach research without a gatekeeper standing in the way. She's also clear about what comes with that. As she put it on our Close the Listening Gap webinar, "where access is given, responsibility is added too."

Guardrails instill that responsibility. Martin Garcia, Staff UX Researcher at NinjaTrader, requires answers to stakeholder questions to draw on primary themes backed by at least three participants, a rule he described on our How Modern Researchers Are Using MCP webinar.

For Mahad Bullo, Senior Staff UX Researcher at Included Health, the goal is infrastructure. "If I weren't around tomorrow, I want us to have the best insights infrastructure of any company in our space. That is what I am building toward," he says. The real test of ownership is what keeps working without you.

Steps to build the context layer

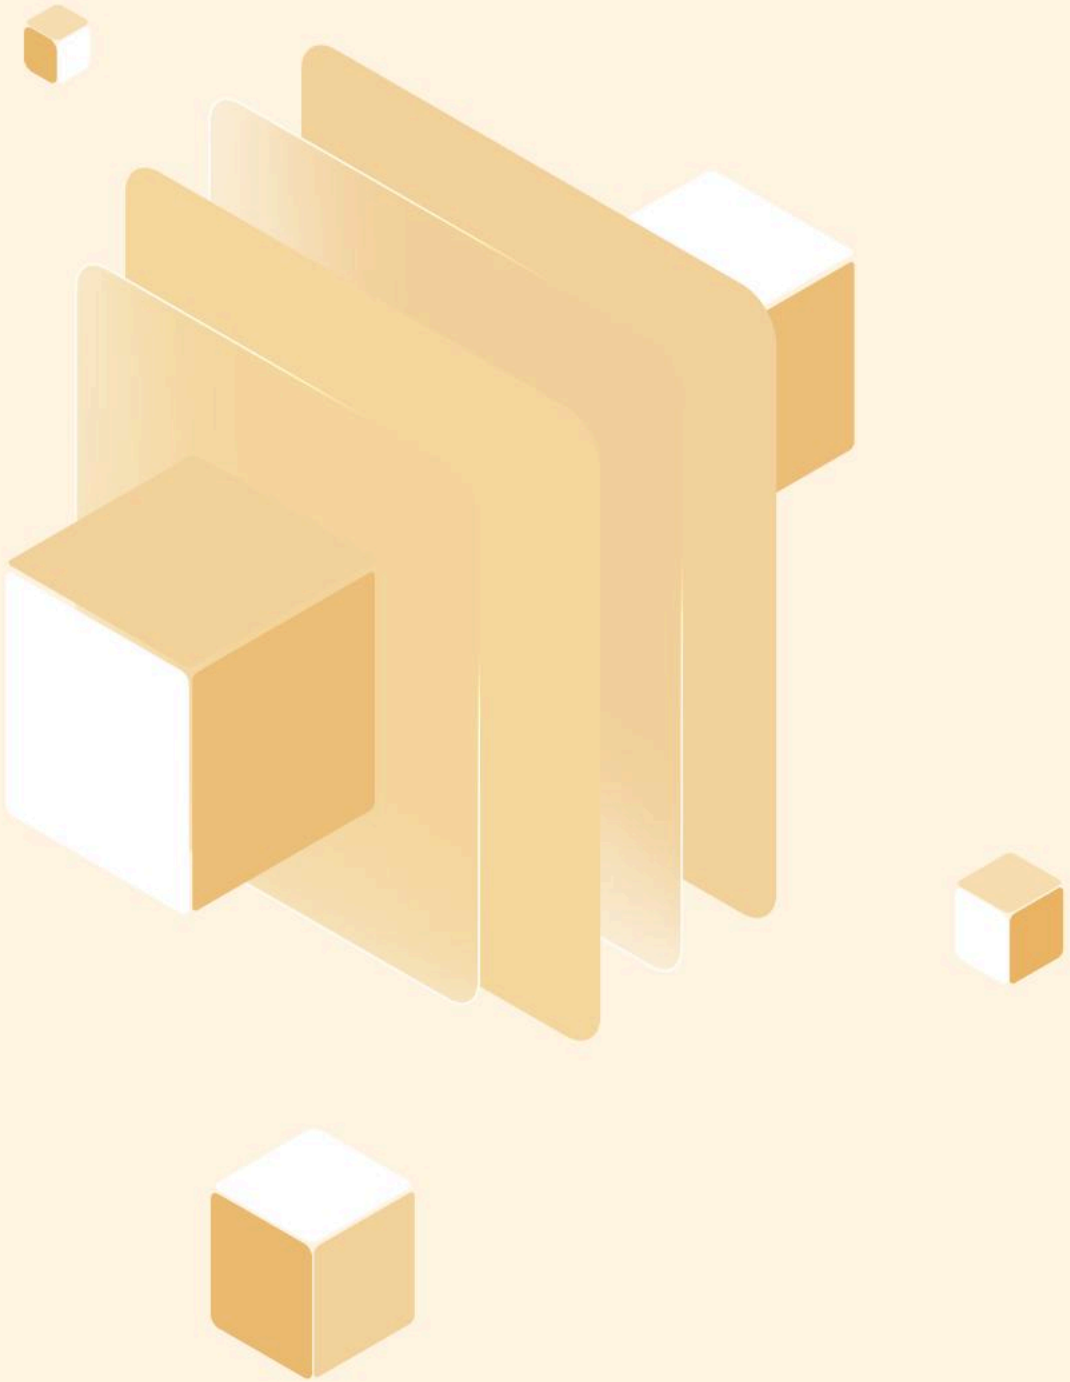

The playbook

Steps to build the context layer

Building a customer context layer comes down to four steps. Most teams aren't starting from scratch, since the research and feedback they already have gives them a foundation to build on.

1. Structure the data for AI

Order matters. Bringing scattered sources into one place without structure makes AI more likely to draw unreliable conclusions from them.

Start with structure: a shared taxonomy, label descriptions that tell an AI what each label means, and project overviews that say what each study set out to learn.

Mahad Bullo builds every project around the AI, as he describes in Building Research Influence Through Participation. He sets up tags before the data arrives and writes label descriptions as cues for the AI to follow.

Then bring the sources in and keep each one linked to where it came from. Connect the people behind them, too. The layer should recognize one customer across an interview, a survey, and a ticket, and never mistake an interviewer's question for a customer quote.

2. Give the layer your company's truth

Customer context tells AI what customers think, say, and need. Company context tells AI what your business has already settled as true, such as how the product works, what your team calls things, and which decisions are final. That's your company's truth, and with it, AI can read customer research in the context of the business.

Some teams keep company context in a separate workflow. When we interviewed research teams, some described connecting a general AI assistant to their company wiki for company context, then using a research tool for customer evidence.

It's a reasonable workaround, but it splits context in two, so no single answer draws on both: what customers said and what the business has decided.

Bringing company truth into the layer is a team effort, and much of it already exists, even if scattered. Product knows the product facts and which decisions are settled. Marketing knows the brand guidelines and how the company describes itself. Legal and security know what must never leave the layer and which permissions and redaction rules keep it protected. Once those teams agree on what's true, the layer can treat it as the final word in every answer.

3. Connect it to every AI tool your teams use

A layer gets used when people can reach it from the tools they already work in, whether that's an AI assistant, a design tool, Slack, or an internal agent. For AI tools, MCP is the most direct route. Connect the layer once, and every compatible assistant, agent, and design tool draws on the same customer knowledge.

The design team at athenahealth connected their research to Figma Make, the AI tool where they build working prototypes. "By using MCP connectors, we are no longer designing in a vacuum. We are designing with the context of our user research right inside our Figma files. This makes the design process less linear because user insight stays present deeper into the evolution of a project," says Katherine McVey, Senior Product Designer.

That matters because people rarely leave the tools they already use. Alexis McNutt Unis, Customer Research Lead at OnePay, sees the same thing from the research side. "A lot of the tooling that I have ... requires a lot of training for folks," she said on our Close the Listening Gap webinar. "MCPs are ... breaking down barriers because I know that everybody's within their [Claude Code] or whichever AI of choice."

Start with one workflow. Map it end to end, list the tools used at each step, and prioritize connecting the layer to those tools.

> **"By using MCP connectors, we are no longer designing in a vacuum. We are designing with the context of our user research right inside our Figma files."**
>
> 
>
> **Katherine McVey**
> Senior Product Designer, athenahealth

4. Keep it current, and prove it's working

This step decides whether the layer lasts. Make sure the layer updates itself as new evidence arrives, and give each ongoing data flow an owner to keep it running.

The healthiest layers are built into how teams ask questions, so checking the layer becomes the first step before any new research. Included Health runs a continuous discovery loop. The team identifies a question, checks what they already know, launches a new study only if needed, and feeds the new insights back in.

Every addition makes a well-built layer more valuable. Each new study, ticket, or call connects to what's already there, so new evidence enriches the old instead of sitting beside it. Over time, that lets teams spot patterns no single study could show, like a frustration that has surfaced quietly for a year or a need that keeps growing across segments.

Those patterns pay off when they shape decisions, so track the decisions the layer informs. At Pantheon, the CEO was hesitant to fund mobile responsiveness. The head of design searched the team's research for "mobile design" and found quotes across 15 studies from the past year, with **12.5%** of users naming it as a pain point. That evidence convinced the CEO.

The framework at a glance

Use this as a quick reference as you work through the steps.

Step	Goal	What good looks like	Who to involve
1. Structure the data for AI	A context layer that AI can read correctly	• Shared taxonomy AI can read • Project overviews • Linked citations	• Researchers/ research ops • Owners of each source (support, sales, CS) • Outside research partners
2. Give the layer your company's truth	Answers you can trust, with governance built in	• Agreed product facts • Company terminology • PII redaction • Permissions that carry over	• Product • Product marketing • Leadership • Legal and security
3. Connect it to every AI tool your teams use	Context people can reach from their tools	• One workflow connected end-to-end, with teams getting cited answers without switching tools	• Owner of each workflow • IT
4. Keep it current, and prove it's working	A living layer that shows its impact	• Automatic updates • An owner for each source • New findings linked to past ones • Tracked decisions	• Owners of each data flow • Research leader • Decision-makers who use it

The hidden cost of building your own context layer

Owning the layer also means deciding whether to build it in-house or buy a solution. Many teams are being asked to build by leveraging the AI tools they already have. But before taking that on, it helps to see the layer for what it really is: ongoing infrastructure. It has to manage everything you've already collected and keep pace with everything you collect next, and that upkeep is the part most teams underestimate.

Upkeep is what keeps the layer current, and that affects whether people keep using it. In our State of Modern Research data, **36%** of respondents with a repository describe theirs as underutilized or stagnant, a common sign that upkeep has fallen behind.

So, the real cost surfaces when it becomes clear what a team has to keep running *after* launch. Ingestion has to stay continuous, and structure, citations, permissions, and governance have to hold up as volume grows. Integrations change, AI models need upgrading, and someone has to support every team that relies on the layer.

That ongoing work also comes with four less-hidden costs:

Tokens

Agents use tokens to build the layer. Then, every question every day costs more, especially when data isn't structured.

Time

Engineering spends months building the layer before the first insight ships. Account for ongoing maintenance as models and connectors change.

Risk

A breach can compromise company knowledge, while downtime leads to research being left out of the decision making.

Opportunity cost

Every sprint spent on internal tooling is a sprint not spent building the products that differentiate your company.

When you use a platform built to run the layer, your team can focus on the tools, agents, and workflows people already use.

◆ In practice

What this looks like with Marvin

Teams across industries use Marvin to build and maintain their customer context layer.

Native integrations and APIs centralize customer and market knowledge, while Organizational Memory maintains company truth. Built-in controls let teams organize their data, redact sensitive information, and manage access for each user.

The result: a shared brain for everything your organization knows about its customers.

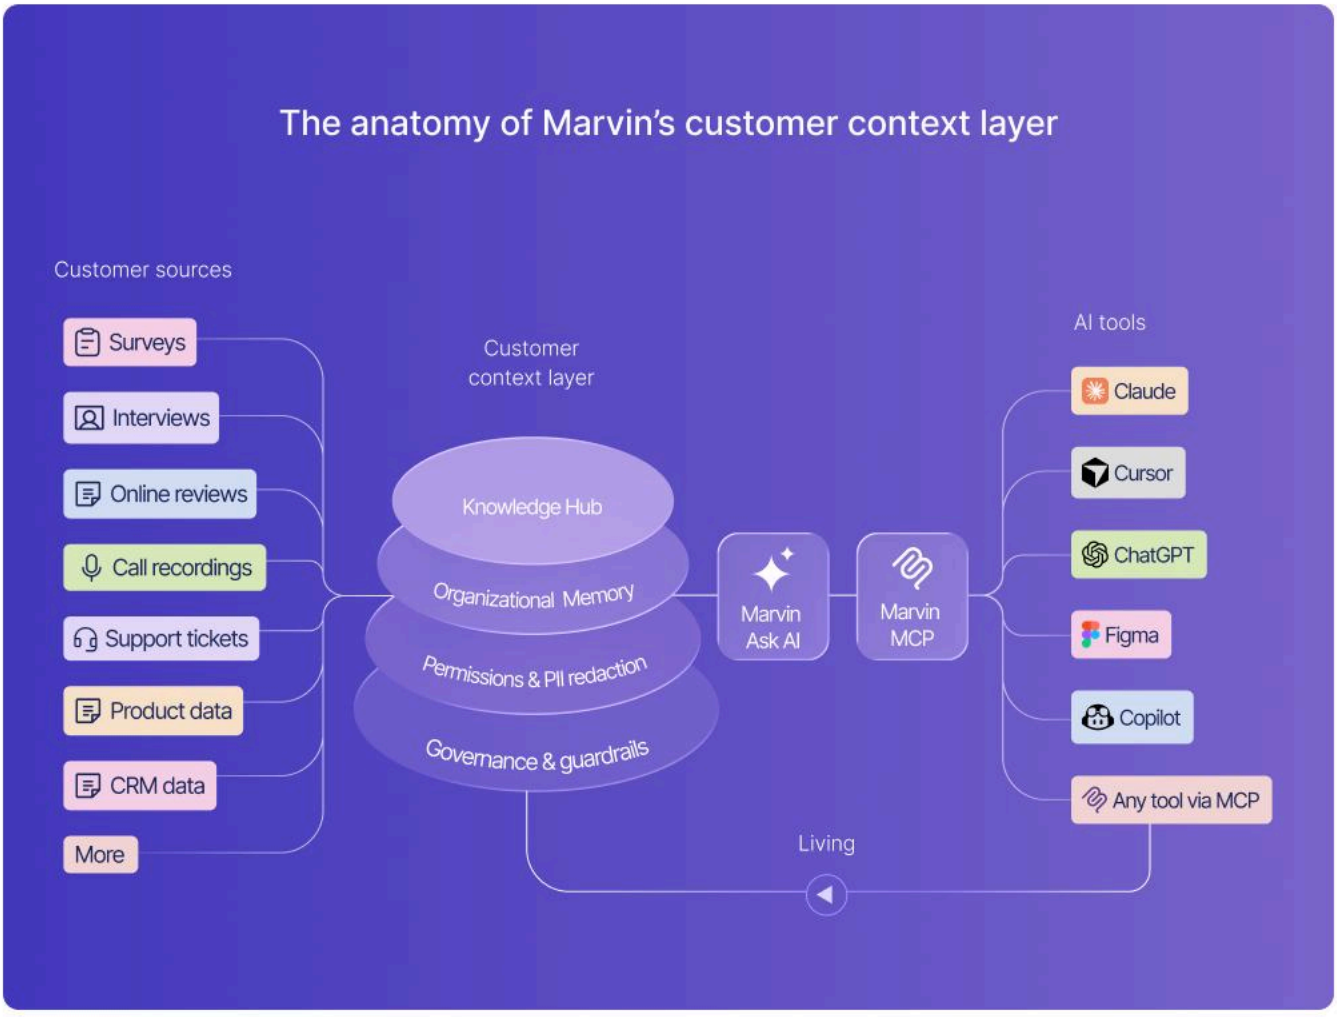

Part of the layer	In Marvin
Complete	**Knowledge Hub and integrations** bring together interviews, surveys, support tickets, sales calls, product feedback, and other customer knowledge in one searchable hub.
Trustworthy	Three things work together: **Ask AI cites** every customer claim to its source, so anyone can check who said it and when. **Organizational Memory** gives AI your company's terminology, product facts, and goals, so answers reflect how your business actually works. **Governance and guardrails** keep your company's data secure. **Automatic PII redaction** helps you stay compliant with data laws. And because permissions carry across tools, Marvin controls who, but also which AI agents only reach what's meant to be shared.
Connected	**Ask AI** makes customer knowledge accessible to teams across the organization. **Marvin MCP** extends that context to compatible AI tools like Claude, ChatGPT, and Copilot, without requiring teams to leave their workflows.
Living	**AI Moderated Interviews and Live Intercept** collect new evidence continuously, and **Mixed Methods Study** captures structured answers and the reasoning behind them in one session. Results land in Marvin already organized alongside everything else in the layer, so they're ready for AI without reprocessing. Integrations keep everyday feedback flowing in too, so the shared context stays current as customers keep talking. **The more knowledge flows in, the richer every answer gets, without anyone rebuilding the layer by hand.**

Bring your whole layer into Marvin, or start with one source and one AI tool, and grow from there. Permissions, citations, and governance work the same at any scale.

4x research output

Included Health quadrupled its research output in a year, with no added researchers or budget.

About Marvin

Marvin is an AI-native customer insights platform that brings together an intelligent knowledge repository, automated customer research tools, and powerful agentic AI in one platform. Teams get a single place to search everything they know about their customers, collect what's missing, analyze it with AI, and get trustworthy answers right where decisions are made. Trusted by research, product, design, and marketing teams at companies like Included Health, Best Buy, and NRG Energy, Marvin helps organizations build the right thing, faster, backed by insights they can trust. For more information, visit *HeyMarvin.com*.

Loved by the world's leading organizations

Microsoft · SIMON KUCHER · vistaprint. · REWE GROUP

Entertainment Partners · HONDA · Lattice · SONOS

MORNINGSTAR · BEST BUY · CRITEO · nrg

See how Marvin can bring your customer knowledge into every AI tool your teams use.

Book a demo

